intel

APPLICATION NOTE

AP-131

March 1982

Intel 2164A 64K Dynamic RAM Device Description

Memory Components
Application Engineering

1. INTRODUCTION

The Intel® 2164A is a high performance, 65,536-word by 1-bit dynamic RAM, fabricated on Intel's advanced HMOS-D III technology. The 2164A also incorporates redundant elements to improve reliability and yield. Packaged in the industry standard 16-pin DIP configuration, the 2164A is designed to operate with a single +5V power supply with ±10% tolerances. Pin 1 is left as a no-connect (N/C) to allow for future system upgrade to 256K devices. The use of a single transistor cell and advanced dynamic RAM circuitry enables the 2164A to achieve high speed at low power dissipation.

The 2164A is the first commercially available dynamic RAM to be manufactured using redundant elements and also features single +5V operation, low input levels allowing −2V overshoot, a wide t_{RCD} timing window, low power dissipation, and pinout compatibility with future system upgrades. These features make the 2164A easy and desirable to use.

2. DEVICE DESCRIPTION

The 2164A is the next generation high density dynamic RAM from the 2118 +5V, 16K RAM. Pin 1 N/C provides for future system upgrade of 64K to 256K sockets. The 2164A pin configuration and logic symbols are shown in Figure 1.

Sixteen bits are required to address each of the 65,536 data bits. This is accomplished by multiplexing the 16-bit address words onto eight address input pins. The two 8-bit address words are latched into the 2164A by the two TTL level clocks: Row Address Strobe ($\overline{RAS}$) and Column Address Strobe ($\overline{CAS}$). Noncritical timing requirements allow the use of the multiplexing technique while maintaining high performance.

Data is stored in a single transistor dynamic storage cell. Refreshing is required for data retention and is accomplished automatically by performing a memory cycle (read, write or refresh) on the 128 combinations of RA_0 through RA_6 (row addresses) during a 2-ms period. Address input A_7 is a "don't care" during refresh cycles.

3. DEVICE OPERATION

3.1 Addressing

A block diagram of the 2164A is shown in Figure 2. The storage cells are divided into four 16,384-bit memory arrays. The arrays are arranged in a 128-row by 128-column matrix. Each array has 128 sense amplifiers connected to folded bit lines.

Figure 3 depicts a bit map of the 2164A and also shows the Boolean equations necessary to enable sequential addressing of the 16 required address bits (A_0–A_{15}). There is no requirement on the user to sequentially address the 2164A; the bit map and Boolean equations are shown for information only.

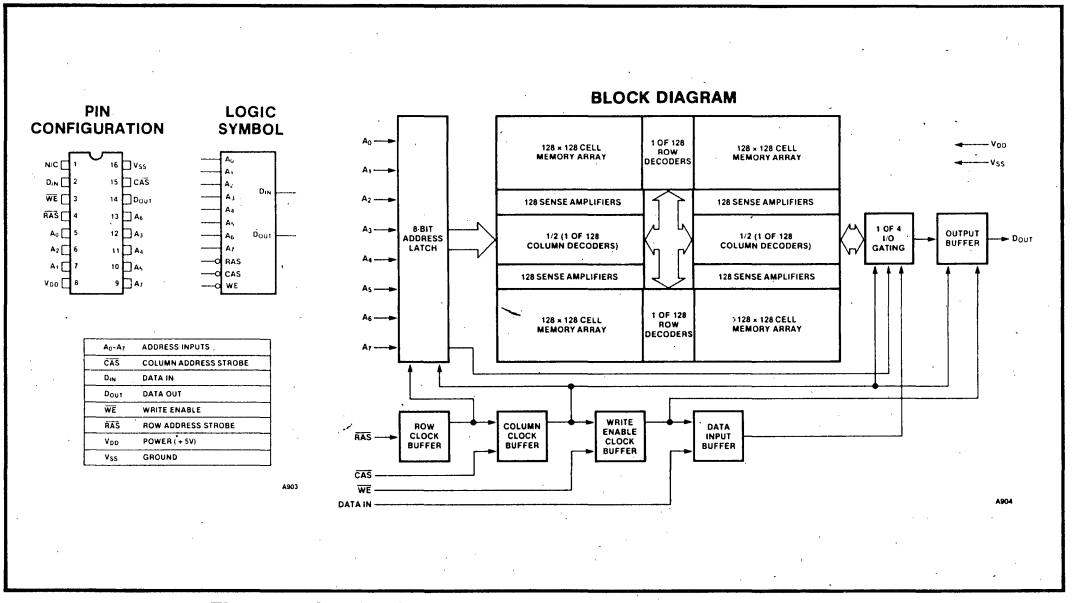

Figures 1 & 2. Intel 2164A Pin Assignments and Block Diagram

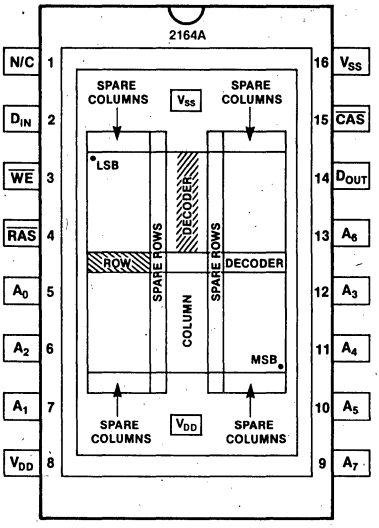

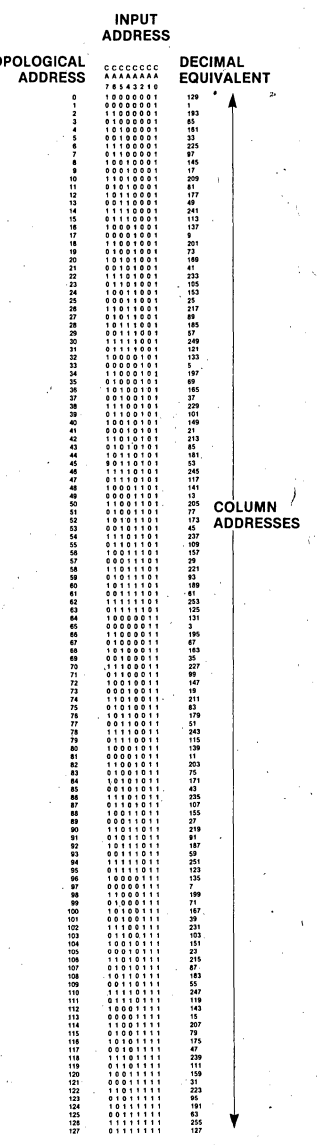

NOTE: Bit Map can be determined from Address Map equations.

2164A

ADDRESS PIN	PROGRAMMED ADDRESS
ROW	ROW ADDRESS SCRAMBLING
RA0	A0R ⊕ A1R
RA1	A1R ⊕ A2R
RA2	A2R ⊕ A7R
RA3	A3R ⊕ A7R
RA4	A4R ⊕ A7R
RA5	$\overline{A5R}$ ⊕ A7R
RA6	$\overline{A6R}$ ⊕ A7R
RA7	A7R
COLUMN	COLUMN ADDRESS SCRAMBLING
CA0	$\overline{AC7}$
CA1	AC6
CA2	AC5
CA3	AC4
CA4	AC3
CA5	AC2
CA6	$\overline{AC1}$
CA7	$\overline{AC0}$
INTERNAL DATA = AOR ⊕ A7C ⊕ D_{IN}	

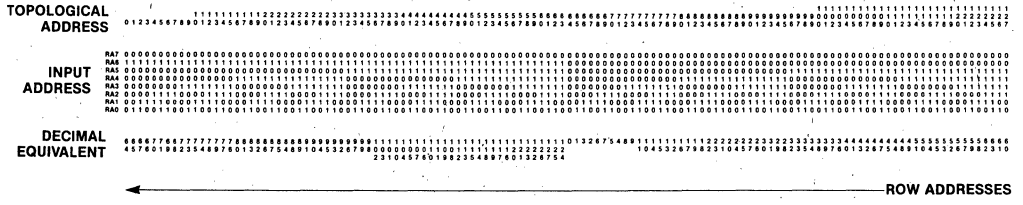

C835

Figure 3. Intel® 2164A Bit Map

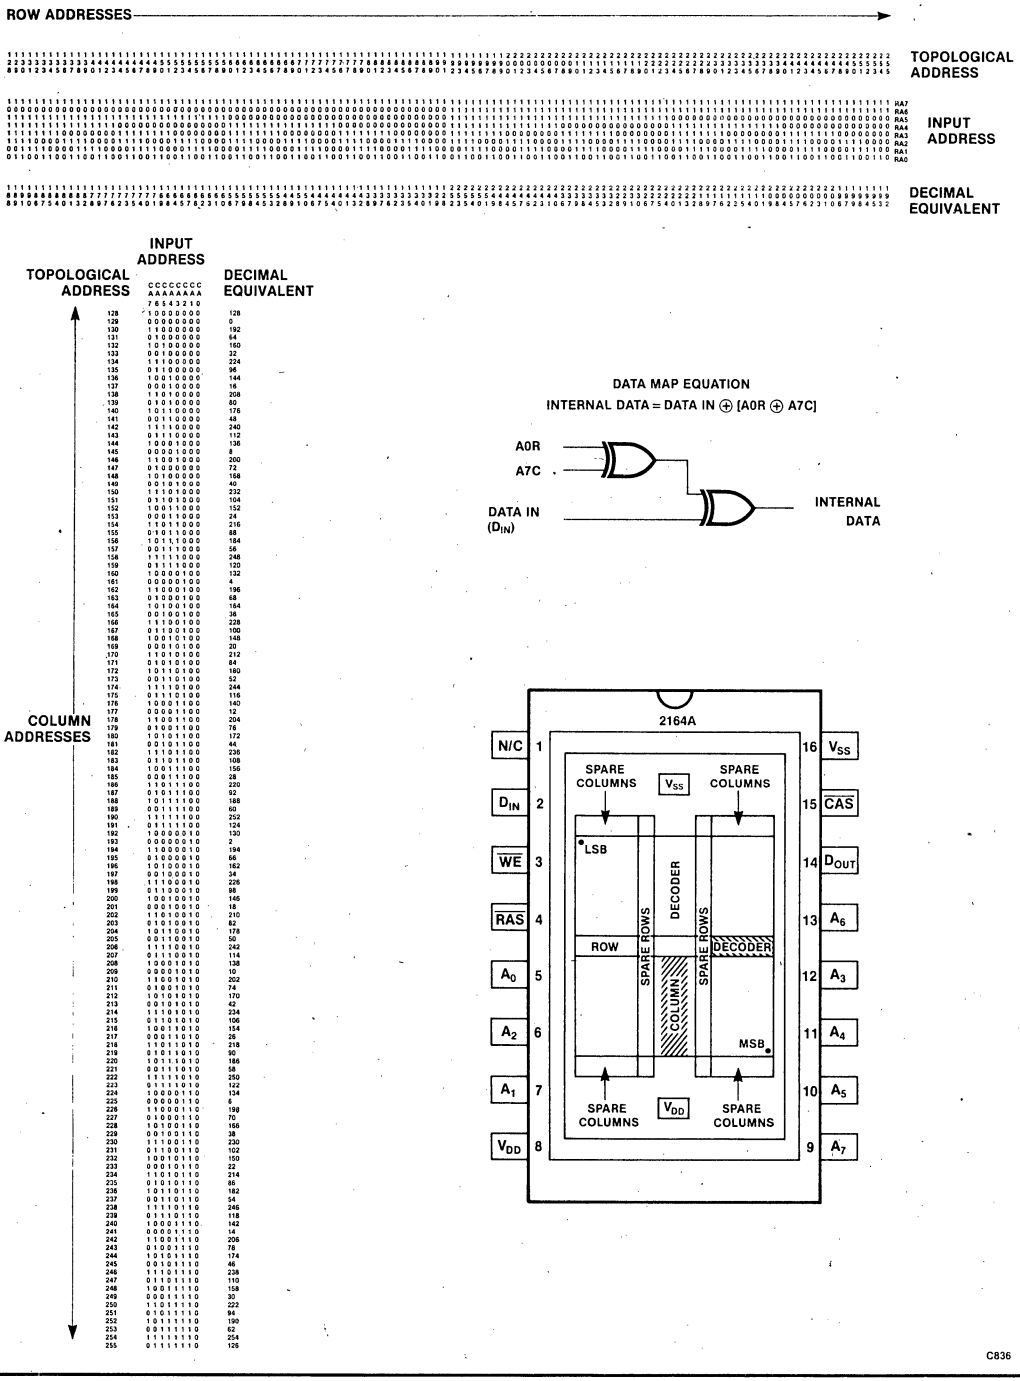

Figure 3. Intel® 2164A Bit Map (continued)

3.2 Active Cycles

When $\overline{RAS}$ is activated, 512 cells are simultaneously sensed. A sense amplifier automatically restores the data. When $\overline{CAS}$ goes active, Column Addresses CA_0–CA_6 choose one of 128 column decoders. CA_7 and RA_7 gate data sensed from the sense amplifiers onto one of the two separate differential I/O lines. One I/O pair is then gated into the Data Out buffer and valid data appears at D_{OUT}.

Because of independent $\overline{RAS}$ and $\overline{CAS}$ circuitry, successive $\overline{CAS}$ data cycles can be implemented for transferring blocks of data to and from memory at the maximum rate — without reapplying the $\overline{RAS}$ clock. This procedure is called Page Mode operation and is described in more detail in Section 4.6. If no $\overline{CAS}$ operation takes place during the active $\overline{RAS}$ cycle, a refresh-only operation occurs: $\overline{RAS}$-only refresh.

3.3 Storage Cell

The basic storage cell is shown in Figure 4. Note that the 2164A uses two dummy cells on each bit line to help compensate for alignment effects. Data is stored in single-transistor dynamic RAM cells. Each cell consists of a single transistor and a storage capacitor. A cell is accessed by the occurrence of row select ($\overline{RAS}$) clocks A_0-A_7 into the address pins, followed by column select ($\overline{CAS}$) multiplexing A_8-A_{15} into the address pins.

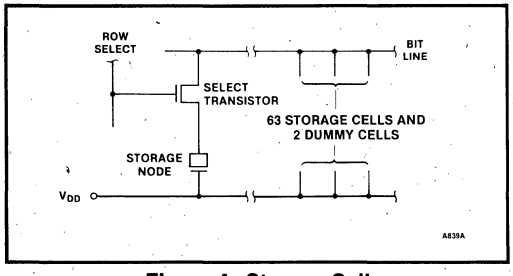

Figure 4. Storage Cell

3.4 Charge Storage in Data Cell

Data is stored in the 2164A memory cells as one of the two discrete voltage levels in the storage capacitor — a high (V_{DD}) and a low (V_{SS}). These levels are sensed by the sense amplifiers and are transmitted to the output buffer. Sensing of stored levels is destructive, so automatic restoration (rewriting or refreshing) must also occur.

The charge storage sensing mechanism for a stored low is described in Figure 5. The V_{DD} storage plate creates a potential well at the storage node. For a stored low, the charge is stored in the cell relative to the storage plate (Figure 5b). The bit sense line is precharged to V_{DD} when $\overline{RAS}$ is high (Figure 5c). During an active cycle, the row select line goes high, and the charge is redistributed (shared) with the bit sense line (Figure 5d). The sense amplifier detects the level from the cell and then reinstates full levels into the data cell via a capacitive bit line restore circuit. At the end of the active cycle, the row select line goes low, trapping the data level charge on the stored cell.

3.5 Data Sensing

The 2164A sense amplifier compares a stored level to a reference level (V_{SS}) in a special, non-addressable storage cell called a dummy cell.

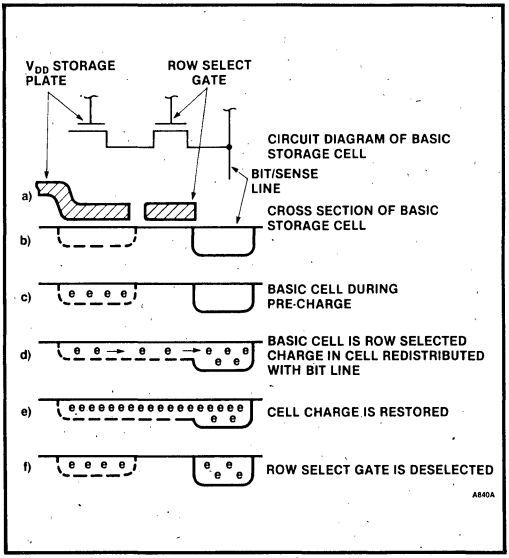

Figure 5. Sensing

Figure 6 depicts a simplified schematic of the 2164A sense amplifier. The sense amp contains a pair of cross-coupled transistors (Q1 and Q2), two isolation transistors (Q3 and Q4), and a common node which goes low with $\overline{SAS}$ (Sense Amp Strobe) and activates the sense amp. The bit-sense lines (BSL and $\overline{BSL}$) run parallel out from the sense amp in a folded bit line approach. Each bit line contains 64 data cells and two dummy cells. The double dummy cell arrangement helps limit the effect of mask alignment on sensing margins by having a dummy cell oriented in the same direction as the data cells.

The folded bit line approach has several advantages, one of which minimizes the effect of interbit line substrate noise and I/O coupling by providing common mode noise rejection. This sense amp arrangement uses metal bit lines and polysilicon word lines.

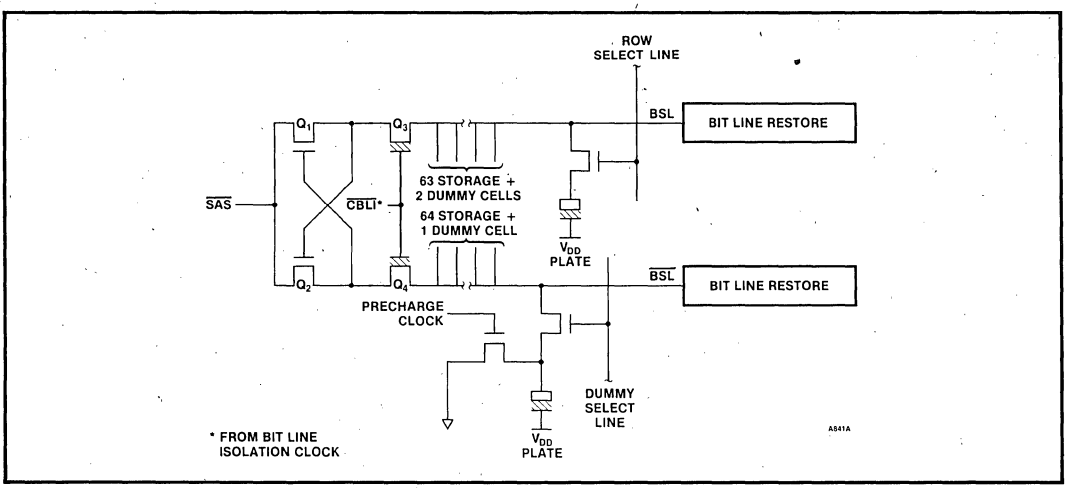

Figure 6. Sense Amp

To eliminate sensing problems, a three-step sensing (Figure 7) is employed in the generation of Sense Amp Strobe clock ($\overline{SAS}$). Device A is triggered by the sense strobe clock. This device pulls down slowly and when fed back, triggers the two gates D and E. When $\overline{SAS}$ is low enough, device B turns on, pulling the $\overline{SAS}$ line lower and at a later time, device C pulls $\overline{SAS}$ down hard. If sensing occurs too quickly, the sense amp becomes sensitive to capacitive imbalance and sensing errors might happen. This design eliminates excessively fast sensing which can occur when two sense strobe clocks are being used.

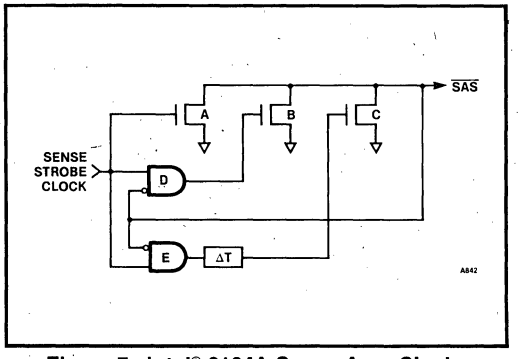

Figure 7. Intel® 2164A Sense Amp Clocks

3.6 Precharge

A precharge period is required after any active cycle to ready the memory device for the next cycle. This occurs while $\overline{RAS}$ is high. The bit lines are precharged to V_{DD}, while the dummy cells are precharged to V_{SS}. During precharge, the row select and dummy select lines are at V_{SS}, isolating the cells from the bit lines. When $\overline{RAS}$ goes low, the precharge clock goes low, ending the precharge period.

3.7 Data Sensing Operation

The row select and dummy select gating are arranged so the selected data and dummy cells are on alternate bit lines of the sense amp (Figure 6). The row select and dummy select lines go high simultaneously, resulting in concurrent charge redistribution on the bit lines. The relationship between the word select lines and the effect of concurrent charge redistribution on the bit lines is shown in Figure 8. An approximate 250 mV differential results from this charge redistribution.

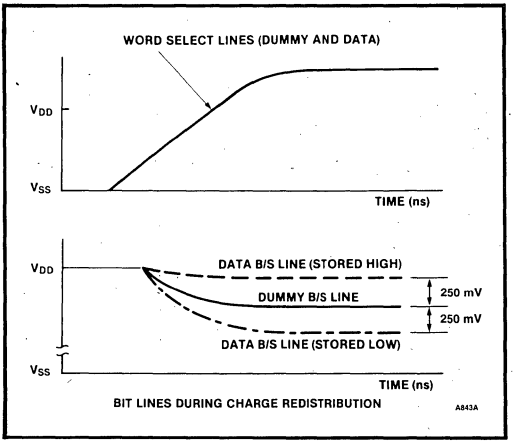

Figure 8. Sensing Voltage Waveforms

After charge redistribution, the sense amp is activated. The sense amp amplifies the differences in the resultant voltages on the bit lines. The line with the lower voltage potential is driven to V_{SS}. The other line remains at a relatively high level, as shown in Figure 9.

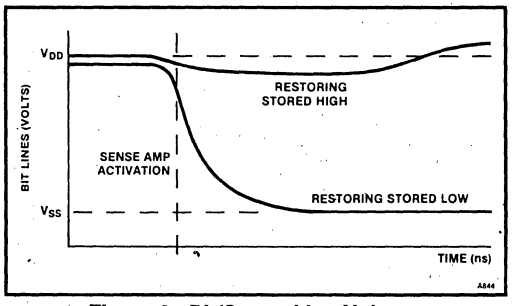

Figure 9. Bit/Sense Line Voltage

The bit line boost circuitry is shown in Figure 10. During sense operations, the boost capacitors are isolated. After sensing, the bit line with a "0" has the capacitor turned off ($V_{GS} \approx 0$) and, conversely, the bit line with a "1" has the capacitor turned on. The boost clock will turn on and boost the 1-level up above V_{DD}, giving maximum charge stored in the cell.

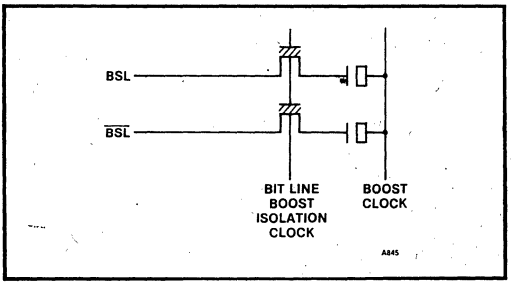

Figure 10. Bit Line Restore

3.8 Data Storage

Figure 11 shows how the I/O lines from each quadrants' sense amps are multiplexed onto the final pair of I/O lines. The I/O is a pair of opposite polarity data lines (I/O and $\overline{I/O}$) which are connected to the Data Input (D_{IN}) and Data Output (D_{OUT}) buffers. Data is differentially placed on the I/O lines during read operation and multiplexed to the final I/O lines. During a write cycle, data is differentially placed on the final I/O lines from D_{IN} and decoded onto the internal I/O lines. Stored levels are determined by CA_7 column and RA_0 row exclusive-ORed product and then exclusive-ORed again with D_{IN} (Figure 3). Stored levels are decoded during D_{OUT} operation and have no effect on device use.

3.9 Address Latches

The 8-bit row and column address words are latched into internal address buffer registers by $\overline{RAS}$ and $\overline{CAS}$. $\overline{RAS}$ strobes in the seven low-order addresses (A_0–A_7) both to select the appropriate data select and dummy select lines and to begin the timing which enables the sense amps. $\overline{CAS}$ strobes in the eight high-order addresses (A_8–A_{15}) to select one of the column decoders and enable I/O operation.

Figure 12 shows a simplified 2164A address buffer. As $\overline{\phi_1}$ goes low, the address input level is trapped via Q1 and similarly, Q2 traps V_{REF}. Since V_{REF} is about halfway between a low (0.8V) and a high (2.4V), either Q3 or Q4 will turn on harder than the other. Then as ϕ_2 becomes active, the cross-coupled latch will change states. As this happens, the load transistor (Q5 or Q6) on the lower side (V_{REF} or A_{IN}) will turn off, minimizing power. As ϕ_3 now becomes active, the address level appears internally at A_X with the complement at $\overline{A_X}$.

The combination of substrate bias and high-speed input buffers allows input overshoots of −2 volts. This is an important specification when designing high-speed switching circuitry driving highly capacitive address busses. Allowing negative overshoots on the address

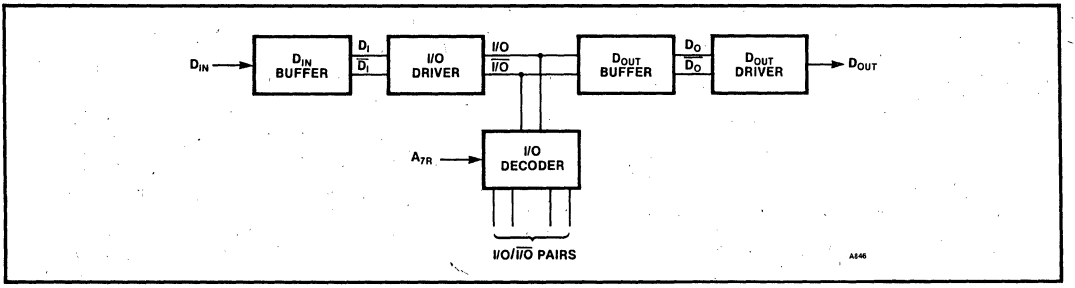

Figure 11. Data I/O

lines means minimum termination of address drivers and increased system performance. This is because a terminated signal (Figure 13) has a slower transition and hence a delay in access time. It is important to note the two advantages to this type of address buffer; first, increased operating speed, and second, a more generous timing window in the multiplexing of the address words.

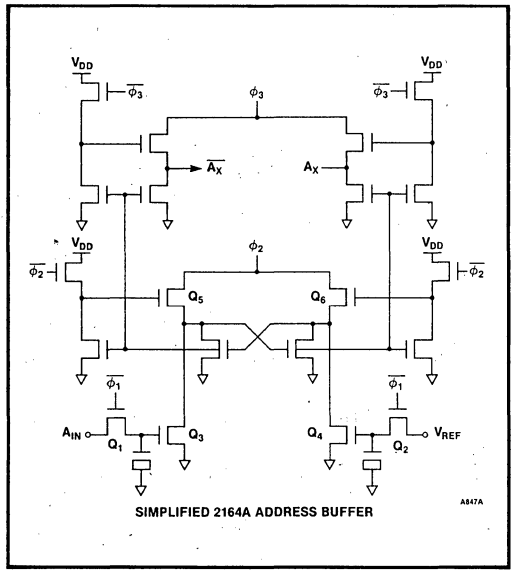

Figure 12. 2164A Simplified Address Buffer Circuitry

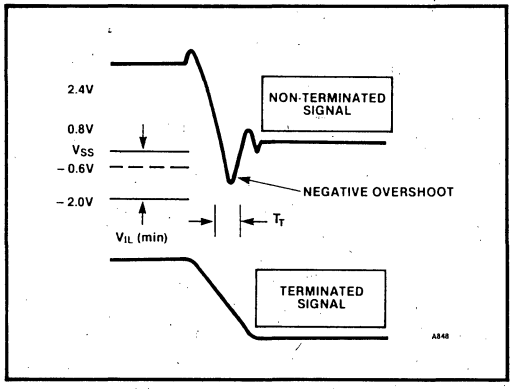

Figure 13. TTL Overshoot

3.10 Data Output Buffer

As shown in Figure 14, the output buffer has a push-pull transistor configuration in which no dc power is dissipated when active.

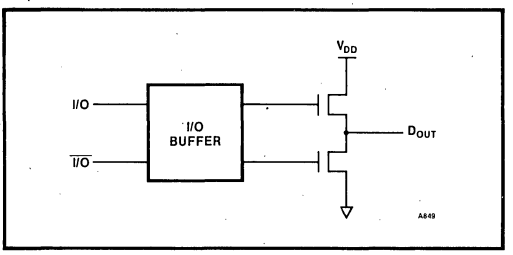

Figure 14. Simplified Output Buffer Circuit

3.11 Data Input/Output Operations

The 2164A contains a Data Input latch which is controlled by the logical NAND function of $\overline{RAS}$, $\overline{CAS}$, and Write Enable ($\overline{WE}$) during the active states (Figure 2). During an early write cycle, where $\overline{WE}$ goes low before $\overline{CAS}$ goes low, the falling edge of $\overline{CAS}$ operates the latch. In a late write (or Read-Modify-Write) cycle, where $\overline{WE}$ goes low after $\overline{CAS}$, the input is latched by the falling edge of $\overline{WE}$.

The 2164A D_{OUT} has three-state capability controlled by $\overline{CAS}$. When $\overline{CAS}$ is at V_{IH}, the output is in a High Impedance (Hi-Z) state. The D_{OUT} states for various operating modes are shown in Table 1. For a Read or Read- Modify-Write cycle, D_{OUT} will remain in the Hi-Z state until the data is valid, whereupon it will go to V_{OH} or V_{OL}, depending on the data.

Table 1. Intel® 2164A Data Output Operation for Various Types of Cycles

Type of Cycle	Data Output State
Read Cycle	Data from Addressed Memory Cell
Early Write Cycle	Hi-Z
$\overline{RAS}$-Only Refresh Cycle	Hi-Z
$\overline{CAS}$-Only Cycle	Hi-Z
Read/Modify/Write Cycle	Data from Addressed Memory Cell
Delayed Write Cycle	Indeterminate
Hidden Refresh Cycle	Data from Addressed Memory Cell
Page Mode Read Cycle (Entry or Internal Cycle)*	Data from Addressed Memory Cell
Page Mode Write Cycle (Entry or Internal Cycle)*	Hi-Z
Page Mode R/M/W Cycle (Entry or Internal Cycle)*	Data from Addressed Memory Cell

* The entry cycle is the first cycle of the page and the internal cycles are the subsequent cycles of the page operation.

For an "Early" Write cycle, D_{OUT} remains in the Hi-Z state which allows "wire-OR" for D_{IN} and D_{OUT}. D_{OUT} is indeterminate for the period between an "Early" Write ($t_{WCS} \geq 0$) and a Read-Modify-Write cycle ($t_{RWD} > t_{RWD}$ min and $t_{CWD} > t_{CWD}$ min). A $\overline{RAS}$-only refresh cycle or a $\overline{CAS}$-only cycle will have no effect on D_{OUT} which will remain in the Hi-Z state. D_{OUT} remains valid from access time until $\overline{CAS}$ goes high. Holding $\overline{CAS}$ low and taking $\overline{RAS}$ high will not affect the state of the D_{OUT}. The D_{OUT} remains valid following a valid Read cycle regardless of the number of subsequent $\overline{RAS}$-only cycles performed on the device up to the t_{CAS} max limit. These secondary $\overline{RAS}$ cycles are $\overline{RAS}$-only refresh cycles to the 2164A.

3.12 Power-On

An initial pause of 500 μs is required after the application of the V_{DD} supply, followed by a minimum of eight (8) initialization cycles (any combination of cycles containing a $\overline{RAS}$ clock such as $\overline{RAS}$-only refresh) prior to normal operation. Eight initialization cycles are required after extended periods of bias (greater than 2 ms) without clocks. The V_{DD} current (I_{DD}) requirement of the 2164A during power on is, however, dependent upon the input levels of $\overline{RAS}$ and $\overline{CAS}$ and the rise time of V_{DD} as shown in Figure 15.

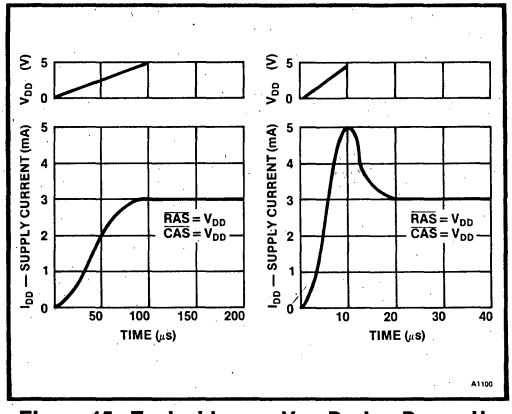

Figure 15. Typical I_{DD} vs. V_{DD} During Power Up

If $\overline{RAS} = V_{SS}$ during power on, the device may go into an active cycle and I_{DD} would show spikes similar to those shown for the $\overline{RAS}/\overline{CAS}$ timings. It is recommended that $\overline{RAS}$ and $\overline{CAS}$ track with V_{DD} during power on or held at a valid V_{IH}.

4. DATA CYCLES/TIMING

A memory cycle begins with a negative transition of $\overline{RAS}$. Both the $\overline{RAS}$ and $\overline{CAS}$ clocks are TTL compatible. The 2164A input buffers convert the TTL level signals to MOS levels inside the device.

$\overline{RAS}$ and $\overline{CAS}$ have minimum pulse widths as specified in the 2164A Data Sheet. These minimum pulse widths and cycle times must be maintained for proper device operation and data integrity. A cycle, once begun, must be within specification.

Figure 16 briefly summarizes the various active cycles which are discussed in paragraphs 4.1 through 4.6.

4.1 Read Cycle

A Read cycle is performed by maintaining $\overline{WE}$ high during a $\overline{RAS}/\overline{CAS}$ operation. The output pin of a selected device remains in a high impedance state until valid data appears at the output within the specified access time.

Device access time, t_{ACC}, is the longer of two calculated intervals:

Eq. (1) $t_{ACC} = t_{RAC}$ or

Eq. (2) $t_{ACC} = t_{RCD} + t_{CAC}$

Access time from $\overline{RAS}$ (t_{RAC}), and access time from $\overline{CAS}$ (t_{CAC}), are device parameters. Row to column address strobe delay time, t_{RCD}, is a system-dependent timing parameter. For example, substituting the device parameters of the 2164A-20 yields:

Eq. (3) $t_{ACC} = t_{RAC} = 200$ ns for 35 ns $\leq t_{RCD} \leq 80$ ns

Eq. (4) $t_{ACC} = t_{RCD} + t_{CAC} = t_{RCD} + 120$ ns for $t_{RCD} > 80$ ns

Note that if 35 ns $\leq t_{RCD} \leq 80$ ns, device access time is determined by equation 3 and is equal to t_{RAC}. If $t_{RCD} > 80$ ns, access time is determined by equation 4. This 45 ns interval (shown in the t_{RCD} inequality in equation 3), in which the falling edge of $\overline{CAS}$ can occur without affecting access time, allows for system timing skew in the generation of $\overline{CAS}$. This allowance for t_{RCD} skew is designed in at the device level to allow the fastest access times to be utilized in practical system designs.

4.2 Write Cycles

4.2.1 EARLY WRITE CYCLE

An early write cycle is performed by bringing $\overline{WE}$ low before $\overline{CAS}$. D_{IN} is written into the selected bit. D_{OUT} remains in the Hi-Z state.

4.2.2 LATE WRITE CYCLE

A late write cycle happens after $\overline{RAS}$ and $\overline{CAS}$ go low. During a late write cycle, t_{RWD} and t_{CWD} ($\overline{RAS}$ and $\overline{CAS}$ delays to Write Enable) minimum timings are not met. Since there is no guarantee that D_{OUT} will remain in a Hi-Z state, the condition of D_{OUT} is indeterminate.

4.3 Read-Modify-Write Cycle (Delayed Write)

A Read-Modify-Write (R-M-W) cycle is performed by bringing $\overline{WE}$ low after $\overline{RAS}$ and $\overline{CAS}$ are low. Here, t_{RWD} and t_{CWD} minimum timings are satisfied. D_{OUT} has had time to become valid and is now latched by $\overline{CAS}$ remaining low. As $\overline{WE}$ goes low, a write begins, transferring the data from D_{IN} to the cell as D_{OUT} remains active with the previous data.

In any type of Write cycle, D_{IN} must be valid at or before the falling edge of $\overline{WE}$ or $\overline{CAS}$, whichever is latest.

4.4 $\overline{CAS}$-Only Cycle

A $\overline{CAS}$-only cycle has no effect on the 2164A. The 2164A remains in the lowest power, standby condition.

4.5 Refresh Cycle

A cycle at each of 128 row addresses will refresh all storage cells. Any memory cycle — Read, Write (Early Write, Delayed Write, R-M-W) or $\overline{RAS}$-only — refreshes the bits selected by the row address combinations of A_0 through A_6. Both 32K halves are refreshed, as the state of A_7 is irrelevant during refresh.

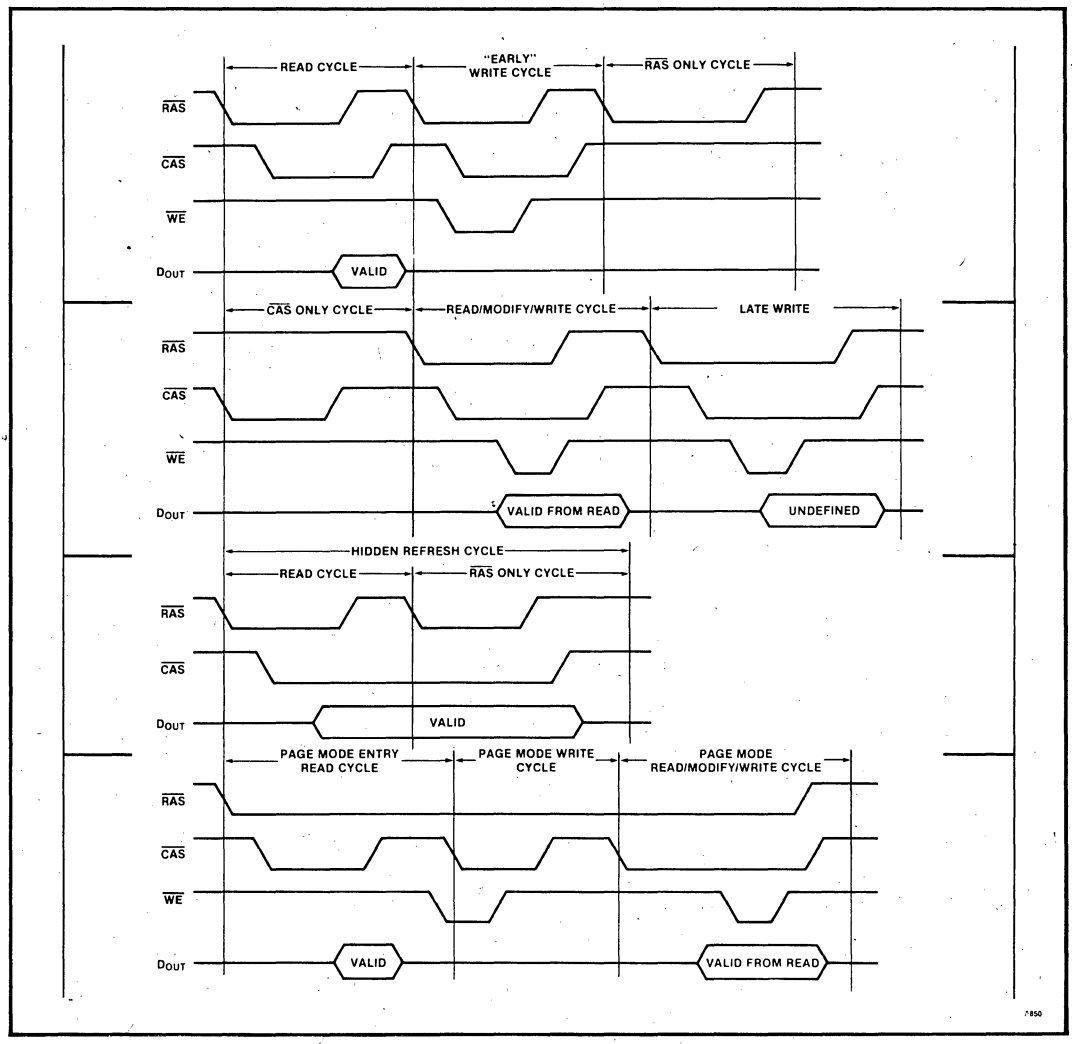

Figure 16. Intel® 2164A Operation of Data Output for Various Active Cycles

4.5.1 READ CYCLE REFRESH

Since A_7 is irrelevant for refresh addressing, a row refreshes 512 cells. The 256 cells in a specific row addressed (A_0–A_6, A_7) are refreshed as are another 256 cells in the row A_0–A_6, $\overline{A_7}$. Therefore, addressing a bit in a row refreshes the 256 cells associated with that row (A_0–A_7). For refresh purposes, row A_0–A_6 and $\overline{A_7}$ is also addressed as another 256 cells. Therefore, successive reads of the 128 row combinations of A_0–A_6 refreshes the entire array of the 2164A.

This refresh mode is useful only when the memory system consists of a single row of devices. When used with more than one row of devices, output bus contention will result.

4.5.2 WRITE CYCLE REFRESH

A Write cycle will perform a refresh. However, the selected cell will be modified to D_{IN}. This may cause a change of state of selected cell, while the other 511 cells are refreshed.

For an Early Write refresh cycle, there will be no output bus contention since the output remains in the Hi-Z state. Bus contention will result for Delayed Write or R-M-W refresh cycles involving more than one row of devices.

4.5.3 $\overline{\text{RAS}}$-ONLY REFRESH

A cycle with $\overline{\text{RAS}}$ active refreshes the 2164A. This is the recommended refresh mode, especially when the memory system consists of multiple rows of memory devices. The D_{OUT}'s may be wired-ORed with no bus contention when $\overline{\text{RAS}}$-only refresh cycles are performed on all rows of devices concurrently. The 2164A D_{OUT} will remain in three-state.

4.5.4 HIDDEN RAS-ONLY REFRESH

The 2164A is designed for "hidden" refresh operation. Hidden refresh accomplishes a refresh cycle following a read cycle without disturbing the D_{OUT}. Once valid, D_{OUT} is controlled solely by $\overline{\text{CAS}}$. After a Read cycle, $\overline{\text{CAS}}$ is held low while $\overline{\text{RAS}}$ goes high for precharge. A $\overline{\text{RAS}}$-only cycle is then performed and D_{OUT} remains valid. However, for operation in this mode, $\overline{\text{CAS}}$ must be decoded along with $\overline{\text{RAS}}$ for the Read and Write cycles. $\overline{\text{CAS}}$ cannot be driven as a common clock to the entire array since it would cause devices being only refreshed to interpret this operation as a $\overline{\text{RAS}}/\overline{\text{CAS}}$ cycle.

4.6 Page Mode Operation

Page Mode operation allows additional columns of the selected device to be accessed at a common row address set. This is done by maintaining $\overline{\text{RAS}}$ low while successive $\overline{\text{CAS}}$ cycles are performed.

Page Mode operation allows a maximum data transfer rate as $\overline{\text{RAS}}$ addresses are maintained internally and do not have to be reapplied. During this operation, Read, Write and R-M-W cycles are possible. Following the entry cycle into Page Mode operation, access is t_{CAC} dependent. The Page Mode cycle is dependent upon $\overline{\text{CAS}}$ pulse width (t_{CAS}) and the $\overline{\text{CAS}}$ precharge period (t_{CPN}).

5. SYSTEM DESIGN CONSIDERATIONS

Calculating total 2164A power consumption is a simple task. To illustrate the method of calculating power, an example system organized as 256K words by 16 bits is assumed.

The first step is to compute the total 2164A current by summing the three individual V_{DD} 2164A supply currents: (1) operating current (I_{DD0}), (2) standby current (I_{DDS}), and (3) refresh current (I_{DDR}). The total 2164A power consumption equals the 2164A current multiplied by the maximum supply voltage (V_{DD}). Total system power consumption is determined by adding the support circuitry power requirements to the total 2164A power.

Examples of these calculations, along with a power/bit determination, are presented in following sections.

5.1 Power Calculations

5.1.1 OPERATING CURRENT (I_{DD0})

Active operating current is determined by the following equation:

Eq. (1) $I_{DD0} = (I_{DD2} + I_{DDLO})K$

Where: I_{DD0} = the operating V_{DD} supply current.

K = the number of active devices (selected at one time by both $\overline{\text{RAS}}$ and $\overline{\text{CAS}}$).

I_{DDLO} = the 2164A output load current (output leakage current plus the load devices input current). For example, if four devices are dot ORed on the output line, the output leakage current is the sum of the input current (I_{IN}) for the load plus the three leakage currents (I_{LO}) for the three devices standby.

5.1.2 STANDBY CURRENT (I_{DDS})

Standby current is determined by the following equation:

Eq. (2) $I_{DDS} = I_{DD1} \times M$

Where: I_{DD1} = the V_{DD} supply current.

M = the number of inactive devices (not selected by $\overline{RAS}$; receiving $\overline{CAS}$-only cycles).

5.1.3 REFRESH CURRENT (I_{DDR})

Refresh current is determined by the following equation:

Eq. (3) $I_{DDR} = (I_{DD3} \times N)\ (t_{RC}/t_{REF})\ (128)$

Where: I_{DD3} = the V_{DD} supply current, $\overline{RAS}$-only cycle.

N = the total number of devices in the system.

t_{RC} = the refresh cycle time.

t_{REF} = the time between refresh cycles.

Since I_{DD3} is not a full-time current, the fraction t_{RC} over t_{REF} represents the duty cycle for one address. There are 128 row addresses active in generating refresh, so the duty cycle is multiplied by 128.

Cycle time has a downward scaling effect on the average operating current according to the following equation:

Eq. (5)

$$I_{DD_{AVE}} = \left[I_{DD2} \times \left(\frac{t_{RC}\ (spec)}{t_{RC}\ (operating)} \right) \right] + \left[I_{DD1} \times \left(1 - \left(\frac{t_{RC}\ (spec)}{t_{RC}\ (operating)} \right) \right) \right]$$

At minimum cycle time, $\dfrac{t_{RC}\ (spec)}{t_{RC}\ (operating)} = 1$,

so that worst case $I_{DD_{AVE}} = I_{DD2}$, but as the cycle time increases, $I_{DD_{AVE}}$ approaches the standby current, becoming 6.3 mA @ 10,000 ns cycle time. Figure 5 in the 2164A data sheet depicts this scaling effect.

5.1.4 TOTAL 2164A POWER

Total 2164A power equals the sum of the three currents multiplied by the worst case supply voltage. This is expressed by the following equation:

Eq. (4) $\text{Power} = (I_{DD0} + I_{DDS} + I_{DDR})\ V_{DD}(\text{max})$

5.1.5 EXAMPLE POWER CALCULATIONS

Assume that we have a 256K word by 16-bit memory system using the 2164A-20 at minimum cycle time. Thus, the following parameters apply:

N = 64 devices in system

K = 16 devices active at one time

M = N-K devices in standby

= 64-16

= 48

Referring to the Intel 2164A Data Sheet[1] and the Intel 8282 Data Sheet[2], we obtain the following values:

I_{DD1} = 5 mA, 2164A-20

I_{DD2} = 45 mA, 2164A-20, t_{RC} = 330 ms

I_{DD3} = 40 mA, 2164A-20, t_{REF} = 2 ms

I_{LO} = 10 μA, 2164A-20

I_{IN} = 200 μA, 8282

To calculate I_{DD0}:

Eq. (1) $I_{DD0} = (I_{DD2} + I_{DDLO})K$

$= (45\ \text{mA} + [3(10\ \mu\text{A}) + 200\ \mu\text{A}])16$

$= 723.68\ \text{mA}$

To calculate I_{DDS}:

Eq. (2) $I_{DDS} = (I_{DD1})M$

$= (5\ \text{mA})48$

$= 240\ \text{mA}$

To calculate I_{DDR}:

Eq. (3) $I_{DDR} = (I_{DD3} \times N)(t_{RC}/t_{REF})(128)$

$= (40\ \text{mA} \times 64)\ \dfrac{330\ \text{ns}}{2\ \text{ms}}\ (128)$

$= (2560\ \text{mA})(.021)$

$= 53.76\ \text{mA}$

To calculate total power:

Eq. (4) $\text{Power} = (I_{DD0} + I_{DDS} + I_{DDR})\ V_{DD}(\text{max})$

$= 5.5\text{V}\ (723.7\ \text{mA} + 240\ \text{mA} + 53.8\ \text{mA})$

$= 5.59\ \text{watts}$

The power/bit is equal to:

Power/Bit = (Total 2164A Power/Number of Devices) (Bits per Device)

$= 5.59(64 \times 65{,}536)$

$= 1.33\ \mu\text{watts/bit}$

5.2 Board Layout

An important consideration in system design is the circuit board layout. A proper layout results in minimum board area while yielding wider power supply and tim-

ing operating margins for increased reliability and easier manufacturability. The key areas of consideration are:

1. Ground (V_{SS}) and power (V_{DD}) gridding
2. Power and ground planes
3. Memory array/control line routing
4. Control logic centralization
5. Power supply decoupling

5.2.1 GROUND AND POWER GRIDDING

Ground and power gridding can contribute to excess noise and voltage drops if not properly structured. An example of an unacceptable method is presented in Figure 17. This type of layout results in accumulated transient noise and voltage drops for the device located at the end of each trace (path).

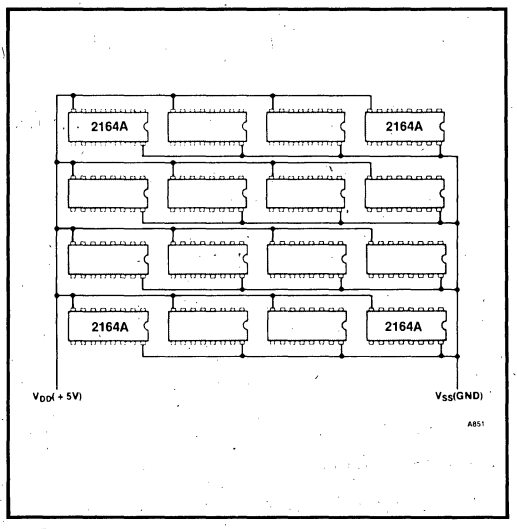

Figure 17. Unacceptable Power Distribution

Transient effects can be minimized by adding extra circuit board traces in parallel to reduce interconnection inductance (Figure 18).

5.2.2 POWER AND GROUND PLANE

A better alternative to power and ground gridding is power and ground planes. Although this requires two additional inner layers to the PC board, noise and supply voltage fluctuations are greatly reduced. If power and ground planes are used, gridding is optional but typically used for increased reliability of power and ground connections and further reduction of electromagnetic noise.

It is preferable on power/ground planes to use circular voids for device pins rather than slotted voids (Figure 19). This provides maximum decoupling and minimum crosstalk between signal traces.

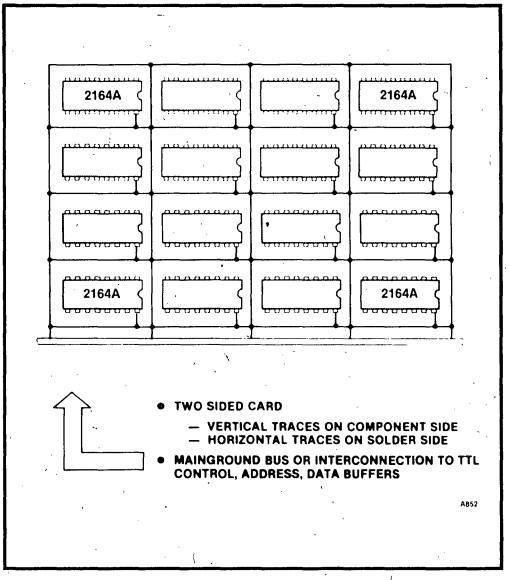

Figure 18. Recommended Power Distribution — Gridding

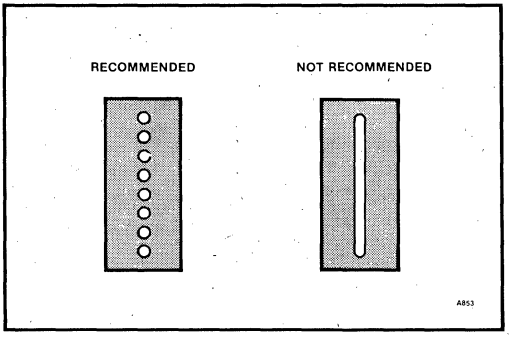

Figure 19. Recommended Voids for Multilayer PC Boards

5.2.3 MEMORY ARRAY/CONTROL LINE ROUTING

Address lines should be kept as short and direct as possible. The lone serpentine line shown in Figure 20 is to be avoided since the devices furthest away from the driver will receive a valid address at a later time than the closer ones. A better way to route address lines is in a comb-like fashion from a central location as shown in Figure 21. Routing control and address signals together from a centralized board area will also minimize skew.

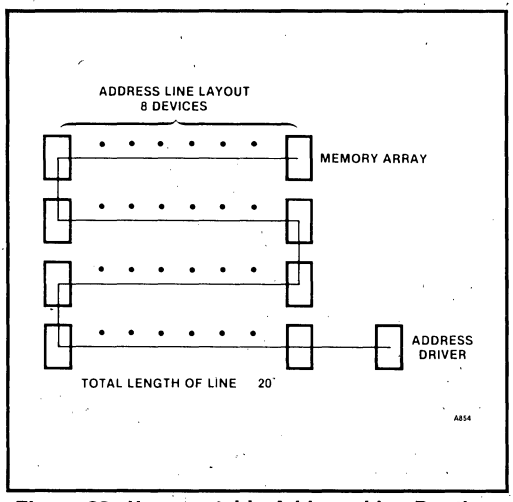

Figure 20. Unacceptable Address Line Routing (Serpentine)

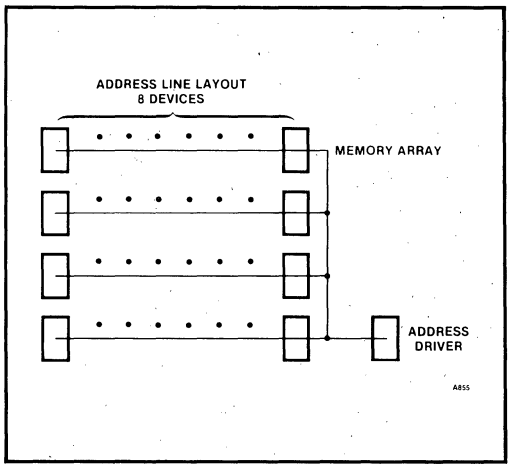

Figure 21. Recommended Address Line Routing

5.2.4 CONTROL LOGIC CENTRALIZATION

Memory control logic should be strategically located in a centralized board position to reduce trace lengths to the memory array. Long trace lines are prone to ringing and capacitive coupling which can cause false triggering of timing circuits. Short lines minimize this condition and also result in less system skew.

A practical memory array layout is shown in Figure 22. Typically, this pattern and its "mirror image" are placed on each side of the memory control logic for a practical memory board design.

5.2.5 POWER SUPPLY DECOUPLING

For best results, decoupling capacitors are placed on the memory array board at each memory location (Figure 22). High frequency 0.1 μF ceramic capacitors are the recommended type, especially for four or more rows of devices. In this arrangement, noise is minimized because of the low impedance across the circuit board traces. Typical V_{DD} noise levels for this arrangement are less than 300 mV.

A large tantalum capacitor (typically one 100 μF per 64 devices) is required at the circuit board edge connector power input pins to recharge the 0.1 μF capacitors between memory cycles.

To calculate decoupling requirements, one considers the current switching of devices from standby to active currents. This involves $I_A = I_{DD2} - I_{DD1}$ (active cycle) and $I_R = I_{DD3} - I_{DD1}$ (refresh cycle). One can then assume some t_B bulk decoupling response time with only one refresh during t_B and minimum cycle time t_C. As a further example, assume only 1/4 of the devices are active at any one time. The amount of charge (Q) requiring decoupling is:

$$Q = I_R t_C + \frac{1}{4} I_A (t_B - t_C).$$

This charge can then be used to calculate the appropriate decoupling capacitance per device. Using Coulomb's law, $Q = CV$, and knowing Q, one picks an acceptable ΔV (< 400 mV) for noise on the V_{DD} lines. The capacitance required is given by $C = Q/\Delta V$. It is important to recognize that C is determined by the current changes in the devices. Minimum cycle time is used for calculating purposes. Lengthening the cycle time will not affect decoupling.

6. THERMAL CHARACTERISTICS

Thermal Characteristics are useful when designing for thermal systems, or for any application where the temperature may go to extremes.

The operating ambient temperature ranges for the 2164A are guaranteed with transverse airflow that exceeds 200 linear feet per minute.

Typical thermal resistance values of the cerdip package at maximum temperature are:

θ_{jA} (@200 fpm air flow) = 47 °C/W
θ_{jc} (still air) = 22 °C/W

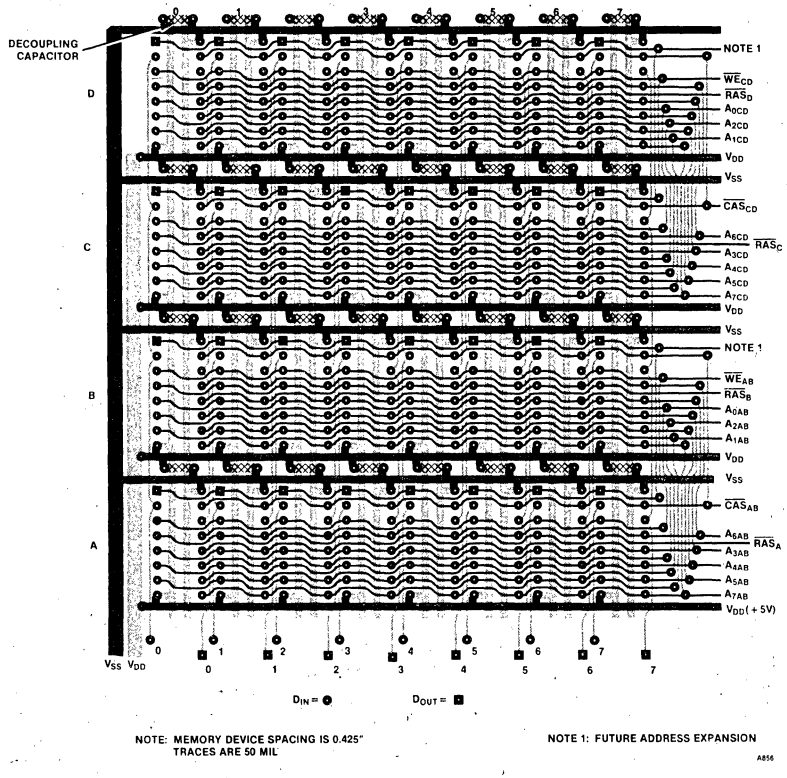

Figure 22. 2164A Memory Array PC Board Layout

7. DESCRIPTION OF REDUNDANT CIRCUITS

The Intel 2164A is the first commercially produced RAM to incorporate redundant elements into the design. Redundancy allows bit-efficient use of silicon by maximizing bits/wafer start. By overstressing and eliminating weak oxide at sort, prior to fusing in redundant elements, long term oxide failures can be greatly reduced. Redundancy makes possible the use of larger die sizes allowing better use of existing fab equipment, and a more conservative layout to utilize larger cell (storage) areas.

In choosing how redundant elements should be organized, single bits, blocks of bits and spare rows and columns were examined. For maximum efficiency, four spare rows and four spare columns were chosen for the 2164A.

The address of a faulty element is programmed into the spare element by electrically opening polysilicon fuses during wafer probe. The basic circuit block diagram for a spare row is shown in Figure 23. The key logic node for the spare row is marked by an (A) on the diagram. When the spare row is not in use, node (A) is held permanently low by transistor (T) whose gate is held high by the spare row enable block. When the spare row is to be used, a fuse is opened within the spare row enable block and the pulldown gate is brought to ground so that the programming elements are enabled. Under control of a fuse, either address true or address complement is transmitted through each programming element. Thus, by blowing the proper fuses, the address of a faulty row in the array is programmed into the spare row.

Figure 24 shows the basic configuration of a programming element. V_G and V_{DP} are special high voltage supplies used only during programming. They are brought on-chip by extra pads probed at wafer sort. These pads are not bonded out to the package but instead, V_G is grounded and V_{DP} is tied to V_{DD} by on-chip transistors. No inadvertent programming can occur at the package level because P_1 cannot turn on and current through the fuse is limited by the transistor connecting V_{DP} and V_{DD}. To blow the fuse, the programming address is brought low, which raises the gate of the programming transistor P to a high voltage. A high current flows through the fuse and it opens. When programming is complete, V_G is brought to ground. If the fuse has been blown, current through depletion transistor D1 pulls node (B) to ground and transfer gate T_2 passes X_i onto X_{pi}. If the fuse has **not** been blown, node (B) stays near V_{DP} and $\overline{X_i}$ is transferred onto X_{pi}.

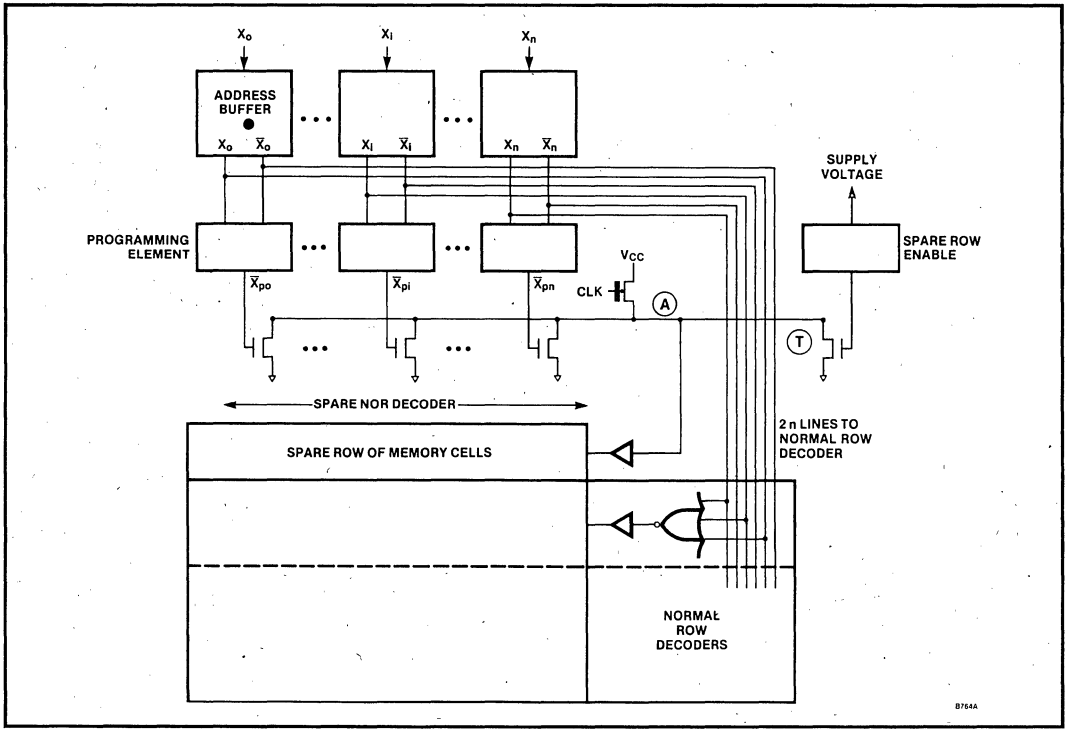

Figure 23. Block Diagram for a Spare Row

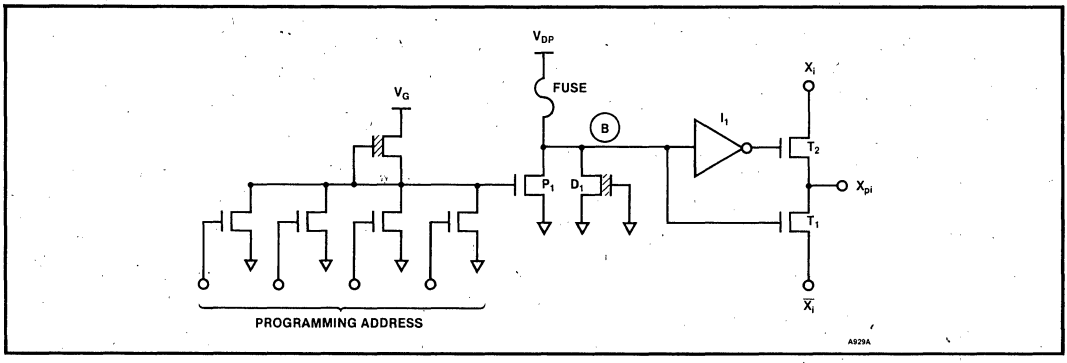

Figure 24. Simplified Circuitry for Programming Element

When the spare row is enabled, one task of the circuit is to deselect the faulty element. Figure 25 illustrates the technique which is used. Whenever any spare select line rises, it causes the "normal element disable" line (NED) to rise as well. NED is connected to one extra input of every normal word select decoder. Thus, when a spare element is selected, it automatically deselects not only the faulty element it replaced, but also every other normal element of the array. The timing of the spare select buffers and the NED generator are optimized to assure that the faulty element is deselected prior to the selection of the spare element.

Another precaution is taken to avoid adverse effects from possible breaks in the faulty select line. If the far end of a broken line were allowed to float, it could present a hazard to data integrity. In the case of a broken word line in the 2164A, word line clamps protect the far end of each row select line from floating high.

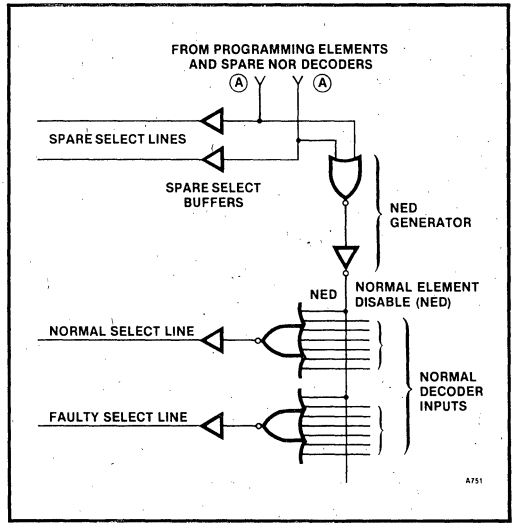

Figure 25. Deselecting a Faulty Element

As mentioned previously, the repair of faulty elements is done during wafer probing. As they come out of fabrication, all spare elements are disabled, allowing full testing of the normal array. Bits are tested not only for hard failures, but also for latent oxide or silicon defects through stressing. The location of any bad bit is stored in the tester's memory. This information is then processed to determine the optimum usage of the spare elements. Then, the spare elements are programmed into their proper logical locations. Finally, the die is tested once more to assure that repair has occurred as planned.

The dice are then assembled as usual. Rigorous class testing is performed to guarantee that the devices meet data sheet specifications in every respect. Both device and system level characterizations have revealed no pattern sensitivity related to the use of redundancy, even when spare elements are intentionally programmed to locations expected to be most susceptible.

Analysis of 2164A devices shows that the worst case patterns do not involve interactions between columns or rows. Replacing the entire row or column introduces no new sensitivity.

The internal delays of redundant element decoding are buried within the internal clocks of the 2164A and have no effect on access time. Figure 26 shows access times for a 2164A before and after repair.

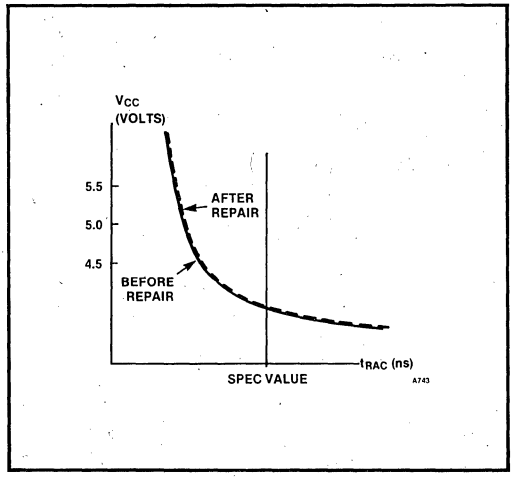

Figure 26. Intel® 2164A t_{RAC} vs V_{CC}

The concept of using redundancy for yield enhancement is well-established. Initially researched by IBM in 1964, Intel has now implemented this concept with the introduction of the 2164A. It is expected that others will follow this lead, and that by the mid-1980's, redundancy will be standard in all memory devices.

8. SUMMARY

The Intel 2164A, made possible by Intel's HMOS-D III technology, introduces a new generation of denser dynamic RAM devices, featuring redundancy, +5V-only TTL-compatible operation, high performance, low power and ease of use. Additional system level design information can be found in Intel Applications Note AP-74, "High Speed Memory System Design Using the 2147H," and AP-133, "Designing Memory Systems For Microprocessors Using the Intel 2164A and 2118 Dynamic Rams."

ADDENDUM

A typical user qualification program of memory devices fits into two categories: device-level qualification and system-level qualification. Occasionally during these programs, failures occur that are not related to the device under evaluation.

At the component level, devices are tested individually for performance to specifications. These tests are usually accomplished with the use of sophisticated software-driven memory testers and environmental handlers. Due to the complexity of the test setup, several problem areas arise. Often testing (software) errors cause failures. Omission of dummy cycles or violation of refresh specifications makes failures invalid. Many times the device under test is remote from the test deck of the system. This can cause excessive power supply noise at the end of the cables. Timing skews, glitches on clock lines and I/O levels at the device are complicated by testing at the end of long cables. Output loading is also critical for the device to perform to specifications.

During system-level qualification, the problems encountered are significantly different. Here the devices are again checked for their performance to specifications. Many devices are simultaneously evaluated whether in a memory system test environment or in an actual system manufactured by the user. Problems can also occur from improper gridding or decoupling on the memory card itself. With the complicated signal paths in a memory system, and the difference between vendor specifications, careful attention must be given to timing and skews not to exceed data sheet values. Errors from timing can result in bus contention or can cause many devices to fail test. Of course, with dynamic RAMs, arbitration between access and refresh modes must be reliable to guarantee the refresh specifications of the RAM.

These problems can be avoided with careful preparation. However, if problems do arise during qualification, don't hesitate to call your local field applications engineer or sales office.

REFERENCES

1. Intel® 2164A Data Sheet, March 1982.